

Identifying lens spaces using discrete logarithms

Greg Kuperberg*
University of California, Davis

We show that if a closed, oriented 3-manifold M is secretly homeomorphic to a lens space $L(n, k)$, then we can compute n and k in randomized polynomial time (in the size of the triangulation of M) with a discrete logarithm oracle. Using Shor's algorithm, a quantum computer can thus identify lens spaces in quantum polynomial time. In addition, k can be computed in functional NP. A given value of k can be certified in randomized polynomial time, specifically in coRP. The idea of the algorithm is to calculate Reidemeister torsion over a prime field that has n th roots of unity.

1. INTRODUCTION

The algorithmic problem of distinguishing or classifying closed d -dimensional manifolds is elementary when $d \leq 2$, provably impossible when $d \geq 4$, and recursive when $d = 3$ [11]. The remaining question is how efficiently we can distinguish closed 3-manifolds; or whether we can distinguish them efficiently with one or another form of help. One small but interesting part of this question is the case of lens spaces. If M is a closed, oriented 3-manifold, conventionally given by a triangulation, then is it a lens space? If so, which one? In this article, we will only address the second question, with the following main result.

Theorem 1.1. *Suppose that M is a closed, oriented 3-manifold given by a triangulation with t tetrahedra, and that secretly $M \cong L(n, k)$. Then n and k can be computed in polynomial time (in t), with the help of an oracle that can compute discrete logarithms in rings of the form \mathbb{Z}/ℓ .*

Recall that each lens space is denoted $L(n, k)$, and is constructed by gluing the top hemisphere of a ball (often imagined as a convex dihedron, a “lens”) to the bottom hemisphere with a rotation of $2\pi k/n$. (It is standard to write $L(p, q)$, but we will write $L(n, k)$ due to the convention in number theory that p denotes a prime number.) The first parameter n is no secret, because if $M \cong L(n, k)$, then we can calculate $H_1(M) \cong \mathbb{Z}/n$ in polynomial time. (The fundamental group $\pi_1(M)$ is also cyclic of order n , but it is less computationally accessible in 3 dimensions; and not at all in higher dimensions.) The second parameter k is more subtle. We can take it to be a prime residue $k \in (\mathbb{Z}/n)^\times$. Reidemeister [17] showed that

$$L(n, k_1) \cong L(n, k_2)$$

as oriented 3-manifolds if and only if $k_1 = k_2$ or $k_1 = 1/k_2$.

In another respect, both parameters are more subtle than one might expect. Suppose that $M \cong L(n, k)$ has t tetrahedra. In the most standard (generalized) triangulation of $L(n, k)$, $n = t$. But there are other families of triangulated manifolds $M \cong L(n, k)$ such that n is exponential in t , and with exponentially many values of k for specific values of n . See Section 3.

One inspiration for our result is a recent result announced by Lackenby and Schleimer [13] to both recognize whether M is a lens space, and if so which one, in the complexity class FNP. (See Section 2 for definitions.) In other words, they provide a deterministic algorithm (a verifier) with the help of a prover who asserts the answer and provides a certificate that it is correct. They have no certificate when M is not a lens space; but there is another NP certificate for this case [9], using the geometrization theorem and assuming the generalized Riemann hypothesis (GRH).

Note that it is easier to answer whether $M \cong L(n, k)$ for any fixed value of n , or when n is written in unary. (So that n rather than $\log n$ is polynomially bounded.) In this case, we can construct the n -fold cover \tilde{M} , and answer $\tilde{M} \stackrel{?}{\cong} S^3$ using the earlier result of Schleimer [20] that 3-sphere recognition is in NP. (It is also in coNP, assuming GRH [9]). We can then compute the Reidemeister torsion of M and calculate k in polynomial time, without discrete logarithms. However, the standard calculation of Reidemeister torsion uses the ring $\mathbb{Z}[\zeta_n]$ or its fraction field $\mathbb{Q}(\zeta_n)$, where ζ_n is an n th root of unity. This is comparable to computing homology in \tilde{M} , and is thus a slow algorithm when n is not controlled.

The idea of our algorithm to prove Theorem 1.1 is simple: We instead compute Reidemeister torsion in a prime field \mathbb{Z}/ℓ with $\ell \equiv 1 \pmod{n}$. It turns out that we need discrete logarithms, but no other artificial help, to interpret the answer. We also need to factor n ; we can do this using discrete logarithms in \mathbb{Z}/n [4].

Our algorithm for Theorem 1.1 has other uses that we summarize in a separate theorem.

Theorem 1.2. *If $M \cong L(n, k)$ is a closed, oriented, triangulated 3-manifold with k secret, then k can be computed in functional BQP and in FNP. If k is provided, then it can be confirmed in coRP. If the prime factorization of n is also provided, then k can be confirmed in ZPP. If the largest prime factor of n is polynomially bounded, then k can be computed in functional ZPP.*

See Section 2 for the definitions of complexity classes used in Theorem 1.2 and in this introduction in general. The gist is that ZPP and coRP are two standards of randomized, polynomial time computation. In coRP, an answer of “no” is certain, while an answer of “yes” is only probably correct. In ZPP, every answer is certain, but the algorithm is only probably fast. Meanwhile BQP is quantum polynomial time; this

* greg@math.ucdavis.edu; Partly supported by NSF grant CCF-1319245

part of Theorem 1.2 just follows from the fact that Shor’s algorithm can compute discrete logarithms and factor integers [21].

To the author’s knowledge, Theorem 1.2 yields the first competitive quantum algorithm for any natural question in 3-manifold topology. For comparison, Aharonov, Jones, and Landau [2] give an algorithm to approximate the Jones polynomial of a knot at a principal root of unity; this algorithm also has a version for 3-manifolds [7]. However, the approximation is usually exponentially poor; any fair approximation that could be useful for geometric topology is #P-hard [12].

This raises the following question: Are quantum computers good at 3-manifold topology more broadly; or do they only understand lens spaces especially well? It is also possible that quantum computers do not even understand lens spaces especially well, if there is a fast classical algorithm to compute k . Note that our algorithm generalizes to distinguish higher-dimensional lens spaces, as explained in the remark at the end of Section 4.2. In this case, a well-known result of Novikov [15] implies that there is no recursive algorithm at all to distinguish lens spaces from other manifolds in dimension $d \geq 5$.

Finally, we note one more corollary of the proof of Theorem 1.1: an upper bound on the work to calculate k without any help.

Corollary 1.3. *If M is a closed 3-manifold given by a triangulation with t tetrahedra where secretly $M \cong L(n, k)$. Then k can be calculated in heuristic time*

$$t^\alpha e^{O((\log n)^{1/3}(\log \log n)^{2/3})}.$$

Corollary 1.3 simply states the running time of the number field sieve algorithm [8] applied to the discrete logarithm problem over \mathbb{Z}/ℓ . Since the work estimate is dominated by the value of n , not t , Corollary 1.3 seems strong in the range

$$(\log n)^{1/3} \gg \log t \gg \log \log n.$$

A more precise work estimate follows from the heuristic prediction of Wagstaff and McCurley [14] that we can take

$$\ell = (2 + o(1))\phi(n)(\log n)^2, \quad (1)$$

rather than just $\ell = O(n)$ as provided by Linnik’s theorem. (Here and in the rest of the article, we use natural logarithm, $\log x = \ln x$, when the base of the logarithm matters.)

ACKNOWLEDGMENTS

The author would like to thank an anonymous MathOverflow user for help with part of the calculation¹.

¹ <http://mathoverflow.net/questions/215852>




Figure 1. Inclusions of some decision complexity classes.

2. COMPLEXITY CLASSES

We briefly review some relevant complexity classes; see the Complexity Zoo [23] for more information.

If Σ is a finite alphabet, then Σ^* denotes the set of all finite strings over Σ , interpreted as data strings. A *decision problem* is a function

$$D : \Sigma^* \rightarrow \{\text{yes}, \text{no}\},$$

while a *function problem* is a function

$$f : \Sigma^* \rightarrow \Sigma^*.$$

A *complexity class* is a set of decision problems or function problems, generally those problems that can be computed with certain resources. It is more standard to start with decision problems, some which are asymmetric between “yes” and “no”. Such a class has a dual class defined by switching “yes” and “no” and denoted with a co prefix; for instance coNP. We are interested in the following decision classes, which are listed together with the type of algorithm that is allowed to compute an algorithm in each of them.

- P – A deterministic algorithm that runs in polynomial time.
- BPP – A polynomial-time randomized algorithm whose answer is probably correct.
- BQP – A polynomial-time quantum algorithm whose answer is probably correct.
- NP – A polynomial-time, deterministic verifier with an omniscient prover who tries to convince the verifier that the answer is “yes”.
- RP – A randomized, polynomial-time prover-verifier whose answer is either a definite “yes” or a probable “no”. Equivalently, the subclass of NP in which a proof certificate can be chosen at random.

- ZPP – A randomized prover-verifier whose answer is certain and who probably succeeds in polynomial time. Equivalently, $\text{RP} \cap \text{coRP}$.

These classes nest as shown in Figure 1. It is conjectured that

$$\text{P} = \text{BPP} \neq \text{BQP} \not\subseteq \text{NP} \neq \text{coNP},$$

while the remaining relation $\text{BQP} \stackrel{?}{\subseteq} \text{NP}$ has no clear evidence for either answer.

Remark. In the standard definition of the probabilistic classes BPP, BQP, RP, and ZPP, the probabilities are bounded away from failure by constants. An algorithm in BPP gives the correct answer with probability at least $2/3$; an algorithm in ZPP finishes quickly with probability at least $1/2$; etc. By using repeated trials, the probabilities can be amplified to be exponentially close to 1. This fact, combined with constructions of cryptographic pseudorandom number generators, is taken as evidence that $\text{P} = \text{BPP}$.

Some decision complexity classes have standard functional counterparts. One convention is to add an F prefix; for instance, FNP. One can also just say “functional NP” or “functional BQP”.

3. HARD LENS SPACES

In this section, we will construct lens spaces $M \cong L(n, k)$ where n is much larger than the number of tetrahedra t , and k is unpredictable. According to Theorem 1.2, identifying M is easy if n is a smooth number (meaning that it has no large prime factors). We cannot prove that our construction can produce values of n with large prime factors, but computer experiments suggests that it can. Also, the specific triangulated manifolds that we construct are easy to identify. However, they can then be obfuscated with a sequence of local moves on triangulations (e.g., Newman-Pacher bistellar moves). Our conclusion is that there are many triangulations of lens spaces that seem difficult to identify.

Proposition 3.1. *There exists a family of triangulated lens spaces $\{M \cong L(n, k)\}$ with $t = t(n, k)$ tetrahedra, such that n is exponential in t and there are exponentially many choices for k for each fixed n .*

Proof. Our construction is equivalent to a well-known construction of lens spaces using Dehn surgery on a chain of unknots [18, Ex. 9H13].

We choose a fixed triangulation σ of the torus $T = S^1 \times S^1$, and we choose two solid tori X_1, X_2 with $\partial X_1, \partial X_2 = T$, and with triangulations σ_1, σ_2 that extend σ . We can describe an element of the mapping class group of T by an element of $\text{GL}(2, \mathbb{Z})$ that describe its action on the homology group $H_1(T)$. For each $1 \leq a \leq 5$, we choose a fixed triangulation τ_a of a torus bundle over an interval, $T \times I$, that connects the triangulation σ of T to itself using the monodromy matrix

$$F_a = \begin{pmatrix} 0 & 1 \\ 1 & a \end{pmatrix}.$$

Our construction is to concatenate a sequence $\{\tau_{a_j}\}_{1 \leq j \leq m}$ of these mapping cylinders together with a solid torus at each end, as in Figure 2. We also assume that $a_1 > 1$. The tetrahedron number t is thus $O(m)$. If the solid tori σ_1 and σ_2 are positioned suitably, then the result is $M \cong L(n, k)$, where n and k are given as a finite continued fraction:

$$\frac{n}{k} = a_m + \frac{1}{a_{m-1} + \frac{1}{\ddots + \frac{1}{a_1}}}.$$

If we let n_j/k_j be the j th partial evaluation, then we can also express the calculation with the recurrence

$$k_j = n_{j-1} \quad n_j = a_j n_{j-1} + k_j = a_j n_{j-1} + n_{j-2}.$$

The answer n/k determines the monodromy numbers $\{a_j\}$ since the continued fraction is unique under the constraint $a_1 > 1$. Since the integers $\{n_j\}$ increase, we obtain the inequality

$$n_j < (a_j + 1)n_{j-1}.$$

If we choose the sequence of monodromy numbers at random, we obtain the probabilistic relation

$$\text{Ex}[\log n_j] < \text{Ex}[\log(a_j + 1)] + \text{Ex}[\log n_{j-1}].$$

Also,

$$\begin{aligned} \text{Ex}[\log(a_j + 1)] &= \frac{(\log 2) + (\log 3) + \cdots + (\log 6)}{5} \\ &< \log 3.73. \end{aligned}$$

By the law of large numbers, most monodromy sequences produce $n < 3.73^m$. On the other hand, there are $4 \cdot 5^{m-1}$ sequences of length m , so by the pigeonhole principle, some value of n must see exponentially many values of k . Any such value of n must also be exponentially large. In any case, for every choice of numbers $\{a_j\}$, $\{n_j\}$ grows at least as fast as the Fibonacci numbers, which also implies that n is exponentially large. \square

4. ALGORITHM AND PROOFS

4.1. Reidemeister torsion

We review Reidemeister torsion [22] from the algorithmic point of view.

Suppose that

$$C_* = \{C_k \xrightarrow{\partial} C_{k-1}\}_{0 \leq k \leq m}$$

is a finite, acyclic chain complex over a field F . (Reidemeister torsion is well defined for a free complex over any commutative ring, but it is easier to discuss algorithms in the field case.) Suppose in addition that each term C_k has a distinguished basis. Since C_* is acyclic and finite, it is isomorphic to a direct sum of complexes of the form

$$0 \longrightarrow F \xrightarrow{\times 1} F \longrightarrow 0.$$




Figure 2. A lens space $L(n, k)$ as solid tori with triangulations σ_1 and σ_2 , connected by twisted bundles $(S^1 \times S^1) \times I$ with triangulations τ_{a_j} .

Define an *adapted basis* A_* for C_* to be one induced by such a decomposition. In other words, if $\alpha_k \in A_k$ is a basis vector, then either $\partial\alpha_k = 0$, or $\partial\alpha_k \in A_{k-1}$ is another basis vector. Then the *Reidemeister torsion* of C_* is

$$\Delta(C_*) \stackrel{\text{def}}{=} (\det A_0)(\det A_1)^{-1}(\det A_2) \cdots (\det A_m)^{\pm 1},$$

where each A_j is also interpreted as the change-of-basis matrix from the distinguished basis to the adapted basis. The following two facts are standard:

1. Every adapted basis yields the same value of $\Delta(C_*)$.
2. Let C_* be the chain complex of a finite CW complex σ with PL attaching maps, possibly with twisted coefficients, and using the cells of σ as its distinguished basis. Then the Reidemeister torsion $\Delta(C_*)$ is invariant under refinement of σ .

The second fact essentially says that Reidemeister torsion is a PL topological invariant. We just have to be careful because the sign of $\Delta(C_*)$ depends on the ordering and orientation of the cells of σ , and ambiguities in the local coefficient system can also make $\Delta(C_*)$ multivalued.

Given C_* as input, we can find an internal basis by induction working from either end, while determinants can be calculated in polynomial time by row reduction. In conclusion, the Reidemeister torsion $\Delta(C_*)$ can be calculated in a polynomial number of arithmetic operations in the field F .

Let M be a closed, oriented rational homology 3-sphere with a triangulation. We first calculate its untwisted Reidemeister torsion with coefficients in $F = \mathbb{Q}$. Using the orientation, we can canonically augment the chain complex $C_*(M; \mathbb{Q})$ at both ends to obtain the acyclic complex

$$Q_* = \left\{ \begin{array}{ccccccc} 0 & \longrightarrow & \mathbb{Q} & \longrightarrow & C_3(M; \mathbb{Q}) & \longrightarrow & C_2(M; \mathbb{Q}) \\ & & & & \longrightarrow & C_1(M; \mathbb{Q}) & \longrightarrow C_0(M; \mathbb{Q}) \longrightarrow \mathbb{Q} \longrightarrow 0 \end{array} \right\}.$$

Then it is standard that

$$\Delta(Q_*) = \pm |H_1(M; \mathbb{Z})|.$$

The sign is not a topological invariant, because the j -simplices of M are unordered and unoriented, so they only provide $C_j(M; \mathbb{Q})$ with an unordered, unsigned basis. We choose an ordering and an orientation of the simplices such that $\Delta(Q_*) > 0$. We can then use the same ordering and orientation

an other Reidemeister torsion calculation on M with twisted coefficients.

Suppose further that

$$H_1(M) = H_1(M; \mathbb{Z}) \cong \mathbb{Z}/n.$$

Then we can compute n in P by computing the Smith normal form of the differential of the simplicial chain complex $C_*(M)$. Thus, as mentioned in Section 1, n is no secret. We also want an explicit simplicial 1-cocycle ω such that $[\omega]$ generates $H^1(M; \mathbb{Z}/n)$; this too can be calculated in P .

Finally suppose that $M \cong L(n, k)$ is a lens space, and let $F = \mathbb{Q}(\zeta)$, where $\zeta \neq 1$ is a formal root of unity whose order m divides n . We can use the cocycle ω to define a twisted coefficient system $\mathbb{Q}(\zeta)_{\zeta^\omega}$ on M , and let

$$R_* \stackrel{\text{def}}{=} C_*(M; \mathbb{Q}(\zeta)_{\zeta^\omega}).$$

Then a calculation shows that

$$\Delta(R_*) = \zeta^c (1 - \zeta^a)(1 - \zeta^b), \quad (2)$$

where

$$\frac{a}{b} = k^{\pm 1} \in \mathbb{Z}/m.$$

This answer is easy to calculate using the standard cellulation of $L(n, k)$ with one cell in each dimension, as follows. For a convenient choice of twisted coefficients, this CW complex yields

$$0 \longrightarrow \mathbb{Q}(\zeta) \xrightarrow{1-\zeta} \mathbb{Q}(\zeta) \xrightarrow{0} \mathbb{Q}(\zeta) \xrightarrow{1-\zeta^k} \mathbb{Q}(\zeta) \longrightarrow 0.$$

Thus,

$$\Delta(R_*) = (1 - \zeta)(1 - \zeta^k).$$

The formula (2) is the same as this one, except generalized to allow unavoidable ambiguities. The factor of ζ^c arises when we change the choice of ω . Other than their ratio, a and b are ambiguous as well, as follows. First, the value of the torsion (2) does not determine the global sign of a and b , only their relative sign, since

$$\zeta^c (1 - \zeta^a)(1 - \zeta^b) = \zeta^{a+b+c} (1 - \zeta^{-a})(1 - \zeta^{-b}).$$

The formula is also symmetric in a and b , so we cannot distinguish k from $1/k$. This stands to reason, since

$$L(n, k) \cong L(n, 1/k).$$

Second, if we replace $\omega \mapsto x\omega$ for some $x \in (\mathbb{Z}/n)^\times$, then this has the same effect as $\zeta \mapsto \zeta^x$, which then multiplies a , b , and c by x .

As mentioned in Section 1, if $n = m$ and n is exponential in the size of the triangulation of M , then the direct calculation of (2) is not in P because of the field $\mathbb{Q}(\zeta)$ is high-dimensional over \mathbb{Q} .

4.2. Number theory

Recall that

$$\phi(n) \stackrel{\text{def}}{=} |(\mathbb{Z}/n)^\times|$$

is the *Euler phi function*.

We seek a prime number $\ell \equiv 1 \pmod{n}$. Dirichlet's theorem says that there are infinitely many such primes and that they eventually have density $1/\phi(n)$ among all primes. We need an effective version of this theorem, a result which was first obtained by Linnik.

Theorem 4.1 (Linnik [10, Ch. 18]). *If $a, b > 0$ are coprime integers, let $\pi(x; a, b)$ be the number of prime numbers $\ell \leq x$ such that $\ell \equiv b \pmod{a}$. There is a universal constant L such that if $x > a^L$, then*

$$\lim_{a \rightarrow \infty} \frac{\phi(a)\pi(x; a, b)}{x/(\log x)} = 1.$$

Less formally: The usual prime number theorem says that the probability that a randomly chosen $\ell \leq x$ is prime is $\approx 1/(\log x)$. Heuristically, we expect these primes to be equally distributed among all $\phi(a)$ prime congruence classes modulo a . Dirichlet's theorem says that this heuristic is correct, while Linnik's theorem establishes a constant L such that the heuristic becomes true when $x > a^L$.

Remark. Lest the unspecified constant L seem discouraging, the Wagstaff-McCurley conjecture (1) implies that it is fast enough to test consecutive values of $\ell \equiv 1 \pmod{n}$ starting with $p+1$. Computer experiments also suggest that if we randomly choose $\ell \equiv 1 \pmod{n}$ in the range $\ell \leq n^2$, then the probability that ℓ is prime is always at least $1/(4 \log n)$ for each $n \geq 2$.

Proof of Theorem 1.1. If we randomly choose an $\ell \equiv 1 \pmod{n}$ with only linearly more digits than n itself, then it is prime with adequate probability. We need to check whether a candidate ℓ is prime. It is now known that primality is in P [1]; an algorithm in ZPP was found previously [3]. Thus, finding ℓ is in ZPP. Actually, we do not even need ℓ to be prime, we only need n th roots of unity in \mathbb{Z}/ℓ . If the factorization of n is known, then we can find such an ℓ with much simpler methods.

Suppose that $\zeta \in (\mathbb{Z}/\ell)^\times$ satisfies $\zeta^n = 1$ and $\zeta^4 \neq 1$. Then we can follow the calculation in Section 4.1, using \mathbb{Z}/ℓ rather than $\mathbb{Q}(\zeta)$. We can construct the twisted coefficient system $(\mathbb{Z}/\ell)_{\zeta^\omega}$. Copying (2), the answer is still

$$f(\zeta) = f_-(\zeta) \stackrel{\text{def}}{=} \Delta(C_*(M; (\mathbb{Z}/\ell)_{\zeta^\omega})) = \zeta^c(1 - \zeta^a)(1 - \zeta^b).$$

However, this value $f(\zeta)$ can be computed in P once we know ℓ and ζ . We now calculate

$$\begin{aligned} f_+(\zeta) &\stackrel{\text{def}}{=} \frac{f_-(\zeta^2)}{f_-(\zeta)} = \zeta^c(1 + \zeta^a)(1 + \zeta^b) \\ g_+(\zeta) &\stackrel{\text{def}}{=} \frac{f_+(\zeta) + f_-(\zeta)}{2} = \zeta^c + \zeta^{a+b+c} \\ g_-(\zeta) &\stackrel{\text{def}}{=} \frac{f_+(\zeta) - f_-(\zeta)}{2} = \zeta^{a+c} + \zeta^{b+c} \\ h(\zeta) &\stackrel{\text{def}}{=} \frac{g_+(\zeta)^2 - g_-(\zeta)^2}{2} = \zeta^{a+b+2c}. \end{aligned}$$

Thus, using the three evaluations $f(\zeta)$, $f(\zeta^2)$, and $f(\zeta^4)$, we can learn the sum $g_+(\zeta)$ and the product $h(\zeta)$ of ζ^c and ζ^{a+b+c} , and the sum $g_-(\zeta)$ and the product $h(\zeta)$ of ζ^{a+c} and ζ^{b+c} . We can thus solve quadratic equations in \mathbb{Z}/ℓ for all four of these numbers. We can then learn the values of an unordered pair $\{\zeta^{sa}, \zeta^{sb}\}$, with $s = \pm 1$, by taking ratios.

If $n \leq 4$, then we can directly calculate the Reidemeister torsion of M ; so suppose that $n > 4$. If ζ has order n , then we can use discrete logarithm to obtain the unordered pair $\{sa, sb\}$, then take their ratio in \mathbb{Z}/n to obtain $k^{\pm 1}$. To find a suitable ζ , we choose $\alpha \in (\mathbb{Z}/\ell)^\times$ at random and then let

$$\zeta \stackrel{\text{def}}{=} \alpha^{(\ell-1)/n}. \quad (3)$$

If we have discrete logarithm, then we also have factoring, so we can factor n and use the test $\zeta^m \stackrel{?}{=} 1$ for each maximal divisor $m|n$ to determine whether ζ has order n in $(\mathbb{Z}/\ell)^\times$. When $n > 12$, it has order n with probability

$$\Pr[\text{ord}(\zeta) = n] = \frac{\phi(n)}{n} > \frac{1}{\pi \log \log n}.$$

(This follows from [19, Thm. 15] when $n \geq 67$.) So we can find an n th root of unity $\zeta \in \mathbb{Z}/\ell$ with adequate probability, and then compute $k^{\pm 1}$ from $f(\zeta)$, $f(\zeta^2)$, and $f(\zeta^4)$. \square

Remark. The calculation in the proof of Theorem 1.1 is much less special than it may seem. Suppose that we can evaluate a linear combination

$$f(\zeta) = c_1 \zeta^{a_1} + c_2 \zeta^{a_2} + \cdots + c_m \zeta^{a_m}$$

for various roots of unity $\zeta^n = 1$, with known coefficients but unknown exponents. Then any m of these evaluations, e.g., $f(\zeta^j)$ for $1 \leq j \leq m$, yield a system of m polynomial equations with m unknowns. It turns out that any system of polynomial equations over \mathbb{Z}/ℓ of fixed size can be solved in ZPP [5, 6]. As a corollary, we can use the same algorithm to identify higher-dimensional lens spaces.

Proof of Theorem 1.2. As mentioned in Section 1, integer factorization and discrete logarithm are both in BQP by Shor's algorithm. Thus, calculating k is immediately in BQP as well. For the FNP result, the prover can provide these data values:

1. The value of k .
2. A prime $\ell \equiv 1 \pmod{n}$.
3. An n th root of unity $\zeta \in \mathbb{Z}/\ell$.
4. The prime factorization of n .

Using this certificate, a deterministic verifier can confirm that ζ has order n , and then follow the calculation in the proof of Theorem 1.1 to calculate ζ^k . (Even though k is ambiguous, the prover will know whether it is to be k or k^{-1} , since the verifier is entirely predictable.) The verifier also has to know to trust the prime factorization of n . For this purpose we can again use that primality is in P, or in $\text{ZPP} \subseteq \text{NP}$; there is also a much simpler earlier construction that primality is in NP due to Pratt [16].

If the verifier only sees the value of k , then we can still search for ℓ at random, test its primality, and then choose $\alpha \in (\mathbb{Z}/\ell)^\times$ at random and exponentiate it to obtain ζ as in (3). If $\zeta^4 = 1$, then we reject it and choose another value of α .

Otherwise it has order $m|n$, with $m = n$ with adequate probability. In this case, we can learn the unordered pair $\{\zeta^{sa}, \zeta^{sb}\}$, and calculate the test $(\zeta^{sa})^k \stackrel{?}{=} \zeta^{sb}$, possibly with a and b switched, and as before with $s = \pm 1$. This gives us the value $k^{\pm 1} \in \mathbb{Z}/m$. Since $m = n$ with adequate probability, and or otherwise since n could be the lcm of several values of m , this shows that $M \stackrel{?}{=} L(n, k)$ is in coRP when k is part of the input.

Finally, if n is polynomially smooth, then it can be factored in polynomial time; a randomly chosen ζ can be confirmed; and there is a simple algorithm to compute discrete logarithms, as follows. Suppose that we know $x = \zeta^a$ and that we know a small factor $m|n$ (not necessarily prime). By induction, we first solve the equation $x^p = (\zeta^m)^a$ to learn $a \in \mathbb{Z}/(n/m)$; more explicitly, we learn

$$a' = a + \frac{nb}{m}$$

with unknown $b \in \mathbb{Z}/m$. We can then solve the equation

$$x\zeta^{-a'} = (\zeta^{n/m})^b$$

for b by exhaustive search. In this version of the algorithm, only finding ℓ and ζ requires ZPP; the rest is in P. \square

-
- [1] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena, *PRIMES is in P*, Ann. of Math. (2) **160** (2004), no. 2, 781–793.
 - [2] Dorit Aharonov, Vaughan Jones, and Zeph Landau, *A polynomial quantum algorithm for approximating the Jones polynomial*, Algorithmica **55** (2009), no. 3, 395–421, arXiv:quant-ph/0511096.
 - [3] A. O. L. Atkin and F. Morain, *Elliptic curves and primality proving*, Math. Comp. **61** (1993), no. 203, 29–68.
 - [4] Eric Bach, *Discrete logarithms and factoring*, Tech. Report CSD-84-186, UC Berkeley, EECS, 1984.
 - [5] E. R. Berlekamp, *Factoring polynomials over large finite fields*, Math. Comp. **24** (1970), 713–735.
 - [6] David G. Cantor and Hans Zassenhaus, *A new algorithm for factoring polynomials over finite fields*, Math. Comp. **36** (1981), no. 154, 587–592.
 - [7] Silvano Garnerone, Annalisa Marzuoli, and Mario Rasetti, *Efficient quantum processing of three-manifold topological invariants*, Adv. Theor. Math. Phys. **13** (2009), no. 6, 1601–1652, arXiv:quant-ph/0703037.
 - [8] Daniel M. Gordon, *Discrete logarithms in $\text{GF}(p)$ using the number field sieve*, SIAM J. Discrete Math. **6** (1993), no. 1, 124–138.
 - [9] Joel Hass and Greg Kuperberg, *3-sphere recognition is in coNP, modulo GRH*, in preparation.
 - [10] Henryk Iwaniec and Emmanuel Kowalski, *Analytic number theory*, American Mathematical Society Colloquium Publications, vol. 53, American Mathematical Society, 2004.
 - [11] Greg Kuperberg, *Algorithmic homeomorphism of 3-manifolds as a corollary of geometrization*, arXiv:arXiv:1508.06720, in preparation.
 - [12] ———, *How hard is it to approximate the Jones polynomial?*, Theory Comput. **11** (2015), 183–219, arXiv:0908.0512.
 - [13] Marc Lackenby and Saul Schleimer, *Lens space recognition is in NP*, Triangulations, vol. 9, Oberwolfach Rep., no. 2, European Mathematical Society, 2012, pp. 1421–1424.
 - [14] Kevin S. McCurley, *The least r -free number in an arithmetic progression*, Trans. Amer. Math. Soc. **293** (1986), no. 2, 467–475.
 - [15] Bjorn Poonen, *Undecidable problems: a sampler*, arXiv:1204.0299.
 - [16] Vaughan R. Pratt, *Every prime has a succinct certificate*, SIAM J. Comput. **4** (1975), no. 3, 214–220.
 - [17] Kurt Reidemeister, *Homotopieringe und linsenräume*, Abh. Math. Sem. Univ. Hamburg **11** (1935), no. 1, 102–109.
 - [18] Dale Rolfsen, *Knots and links*, Mathematics Lecture Series, vol. 7, Publish or Perish, Inc., Wilmington, DE, 1976.
 - [19] J. Barkley Rosser and Lowell Schoenfeld, *Approximate formulas for some functions of prime numbers*, Illinois J. Math. **6** (1962), 64–94.
 - [20] Saul Schleimer, *Sphere recognition lies in NP*, Low-dimensional and symplectic topology, Proc. Sympos. Pure Math., vol. 82, Amer. Math. Soc., 2011, arXiv:math/0407047, pp. 183–213.
 - [21] Peter W. Shor, *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*, SIAM J. Comput. **26** (1997), no. 5, 1484–1509, arXiv:quant-ph/9508027.
 - [22] Vladimir Turaev, *Torsions of 3-dimensional manifolds*, Progress in Mathematics, vol. 208, Birkhäuser, 2002.
 - [23] *The Complexity Zoo*, <http://www.complexityzoo.com/>.